

INTERNET DOCUMENT INFORMATION FORM

A. Report Title: Views on Logistics, Candid Voices, Providing Responsive Logistics Support: Applying LEAN Thinking to Logistics

B. DATE Report Downloaded From the Internet: October 18,1999

**C. Report's Point of Contact: (Name, Organization, Address, Office Symbol, & Ph #): DEPARTMENT OF THE AIR FORCE
AIR FORCE JOURNAL OF LOGISTICS
501 WARD STREET
GUNTER ANNEX
MAXWELL AFB, AL 36114-3236**

D. Currently Applicable Classification Level: Unclassified

E. Distribution Statement A: Approved for Public Release

**F. The foregoing information was compiled and provided by:
DTIC-OCA Initials: __pm__ Preparation Date OCTOBER 18, 1999**

The foregoing information should exactly correspond to the Title, Report Number, and the Date on the accompanying report document. If there are mismatches, or other questions, contact the above OCA Representative for resolution.



Providing Responsive Logistics Support: Applying LEAN Thinking to Logistics

Norman H. Patnode

Let's look at the age-old problem of logistics—to support and sustain the warfighting mission. In order to satisfy this, the logistics system (supply, repair, production, transportation, etc.) must provide customers with what they need when they need it. It must also minimize the cost to the customer.

What is needed to satisfy each of these requirements? In order to provide customers what they need when they need it, the logistics system must repair, produce or purchase things based on forecasted demands. However, in order to minimize the cost to the customer, the logistics system must repair, produce or purchase only what the customers requests and only when they ask for it.

Obviously we have a conflict—we cannot execute (repair, produce or purchase) based on forecasted customer demand and execute based on actual customer demand. The problem is real and has existed since time immemorial. The question before us now is whether we can find a solution to the problem. Let's look at the conflict as it is diagrammed in Figure 1.

A paradox exists when there is a conflict with no apparent solution. Both sides seem to have a logical position but seemingly opposite conclusions. Typically, we compromise in this situation. We repair, produce or purchase things based on a forecast. When the forecast falls short, we frantically attempt to satisfy the urgent backorders. As a result,

customers do not get what they need when they need it, and we never completely minimize the costs to our customers.

A better approach is to challenge the assumptions behind the arrows in Figure 1—assumptions that bind the entities together to form the conflict. If we can remove an arrow by invalidating an underlying assumption, the conflict evaporates.

Using Figure 2, let's examine some assumptions. We get to the assumptions by asking why the tail of the arrow is necessary in order to have the tip of the arrow.

Since it is unlikely that we can eliminate the variability in the logistics process (repair, produce or purchase) or even the uncertainty in the customers' demands, we must continue to forecast or find some way to protect our customers from the effects of the variability.

What if we always had enough stuff on hand to send one to the customers whenever requested? If we could do that, we could stop executing to our forecast. We could simply make replacements through repairing, producing or purchasing. There would no longer be a conflict—the customers would get what they need when they need it, and we would not spend a penny on stuff the customers did not need *now*. This could work. But how do we make sure we actually have enough stuff to protect our customers from the variability, and what happens if we do not?

Let's take the second question first. Take another look at Figure 2. As long as the uncertainty in the customers' demands and the variability in the repair, production or purchasing process continue to impact the customers, the need for forecast-based execution continues to exist. So if

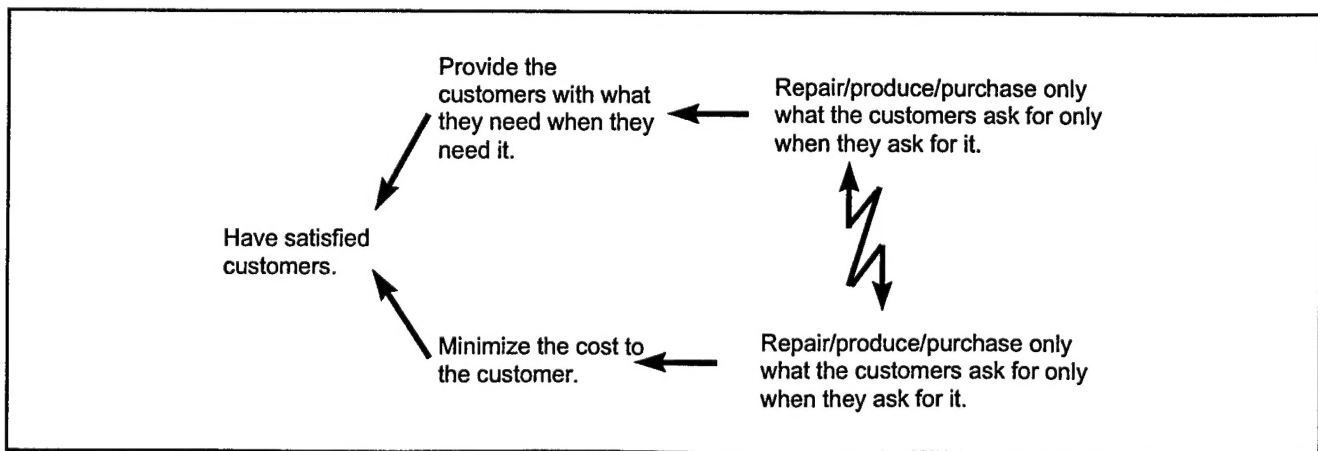


Figure 1. The Logistician's Paradox

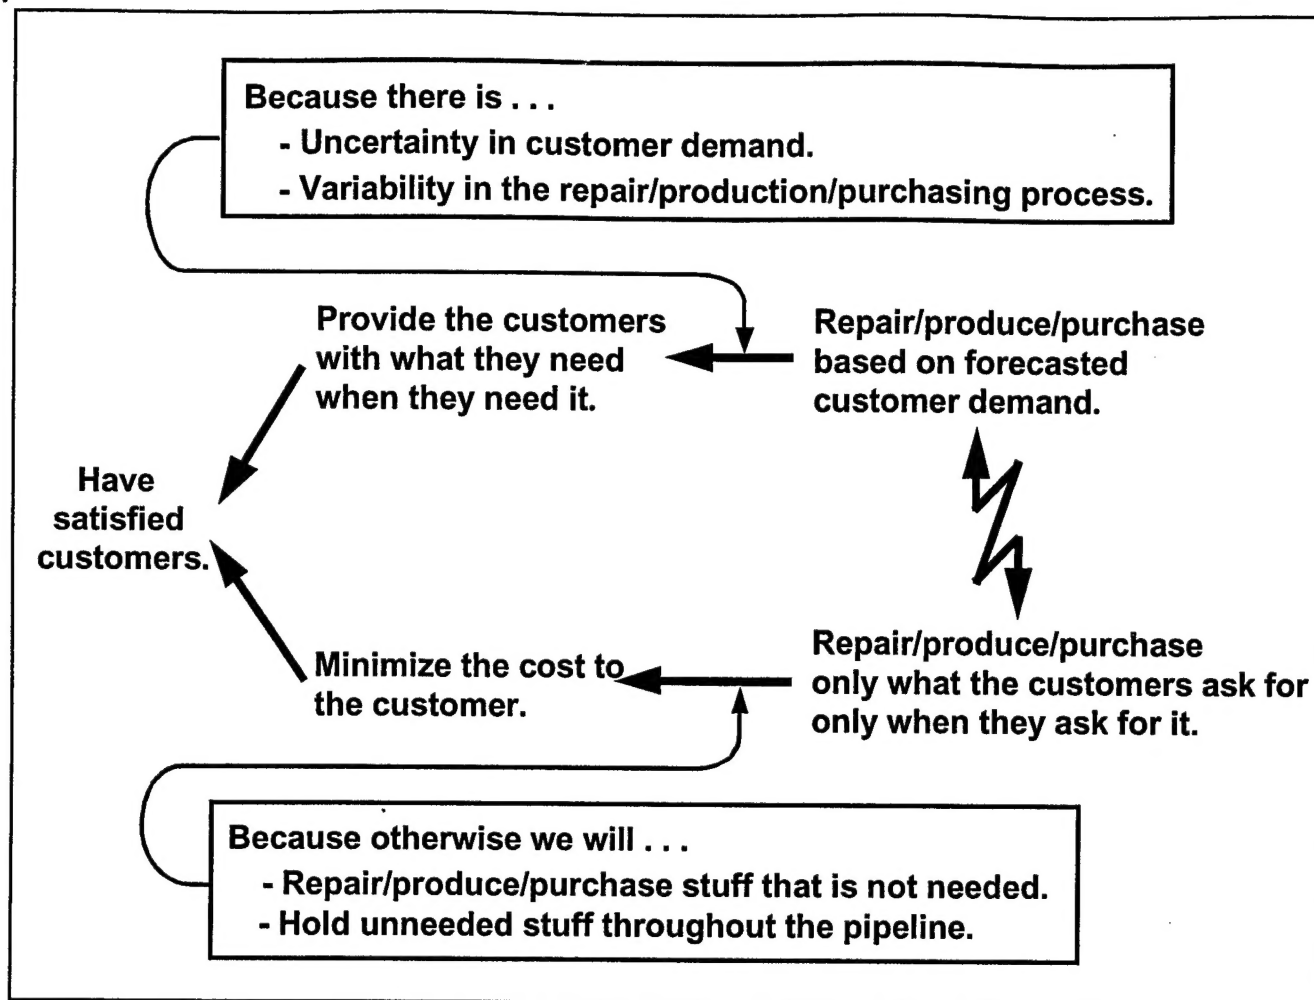


Figure 2. The Underlying Assumptions

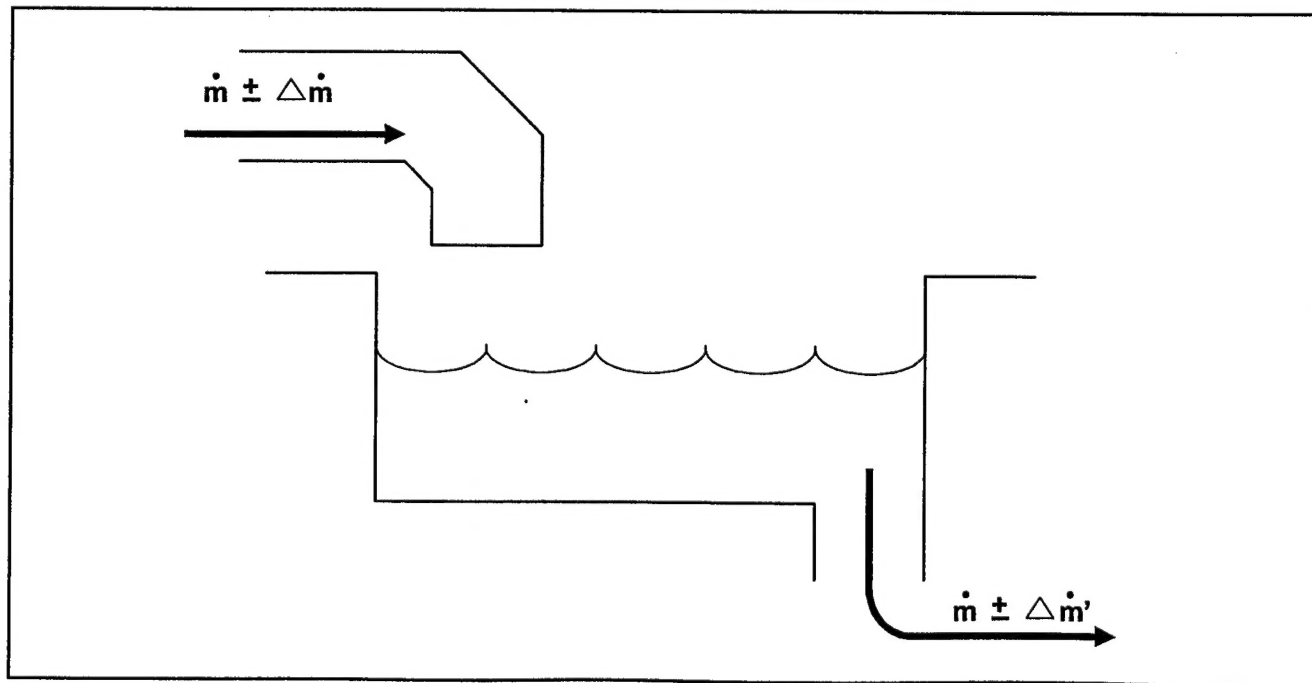


Figure 3. The Protection Buffer

we fail to adequately protect the customers from the variability in the logistics process and we attempt to repair, produce or purchase based on actual real-time customer demands, the customers will not get what they need when they need it. They will always have to wait the full lead time. In short, we will fail to meet the customers' needs.

If we embrace demand-based execution but fail to ensure customers are adequately protected from the variability in the logistics process, we are guaranteed to fail as logisticians.

So let's go back and tackle the first question: How do we make sure we actually have enough stuff to protect the customers from the uncertainty in their demands and the variability in the logistics process?

The first step is to quantify the variability. Then we must decide how much protection is needed. Do we really need to protect against 100 percent of the variability, or is 80 percent good enough? It is a management function to weigh the *cost of protection* against the *impact on the organization's goals*. Finally, we need to continually check to see if the level is providing the needed protection and adjust it as necessary.

Imagine a swimming pool with a pipe feeding in water and a pipe draining water. The input pipe has some average rate of flow with some amount of variability around that average. The output pipe has the same average rate of flow but also has some amount of variability around that average. The variability in the output flow rate is not necessarily the same as the variability in the input flow rate. This is shown in Figure 3. This is your protection buffer. The challenge is to keep enough—and only enough—water in the swimming pool to keep the output pipe from sucking air. Let's take a look at how this can be done. Recall the first step is to quantify the variability as shown in Figure 4.

By looking at the serviceable on-hand balance (the water level in our swimming pool) for our 25-day period, we see that the smallest balance is minus nine assets. (In other words, there are nine unsatisfied customer requirements.) In our example, we said we do not ever want the output pipe to suck air, so we set our level to protect against *all* of the threatening variability. The initial level in our swimming pool needs to be nine assets. (If we choose to protect against only 89 percent of the variability, then we would set the level to be eight assets.)

Remember the final step: continually check to see if the level is providing the needed protection and adjust it as necessary. This is easy, too. Just continue to measure the serviceable on-hand balance and, on a regular basis, set the new level by subtracting the smallest serviceable on-hand balance from the old level. Or in equation form:

$$\left(\begin{array}{c} \text{Percent} \\ \text{New Level} \\ \text{Required} \end{array} \right) = (\text{Protection}) \times \left\{ \text{Old Level} - \left(\begin{array}{c} \text{Smallest Serviceable} \\ \text{On-Hand Balance} \end{array} \right) \right\}$$

Input Pipe (Receipts)	2	0	3	0	0	0	1	0	3	0	0	0	0	0	0	4	0	1	0	1	1	0	5	0
Output Pipe (Requests)	0	0	0	0	0	0	0	9	0	0	0	0	0	0	9	0	0	0	1	0	1	1	0	9
Serviceable Assets On-Hand	2	2	5	5	5	5	6	-3	0	0	0	0	0	0	-9	-5	-5	-4	-5	-4	-4	-5	0	-9

Figure 4. Quantifying the Variability

Note that in order to keep the right amount of water in the swimming pool, it is not necessary to know how long either of the pipes are or what connected them.

The same holds true for any logistics process. By simply watching the serviceable on-hand balance in our protection buffers, we can make sure they provide the required level of protection to our customers.

If our customers are adequately protected from the variability, then we can remove the arrow in the Logician's Paradox, which requires us to forecast. From Figure 5, we see that our solution eliminates the conflict and creates the opportunity for us to truly minimize our cost to the customers by only repairing, producing or purchasing what they request when they ask for it.

But are we missing something? What if we do not have enough assets to fill our protection buffers to the levels needed to protect the customer? Would not that mean, despite all these great ideas, we are still stuck in the conflict. Are we still in a paradox?

In short, yes. But there is a way to get around this dilemma. Let's take a look at something called the Square Root Law, which we can use to increase the protection to our customers, without increasing our requirement for assets.

Today, many of us use forecasts to determine how many assets we need at each location of our distribution system. The Square Root Law makes use of the statistical fact that the accuracy of a forecast increases as you increase the size of the forecast population. (That is why group health and life insurance is so much cheaper than individual policies. As you aggregate the individual forecasts into a group forecast, the forecast becomes more accurate.) As shown, the Square Root Law defines how much better the forecast gets as you aggregate the individual forecasts.

$$\left(\begin{array}{c} \text{Accuracy of the} \\ \text{Aggregate Forecast} \end{array} \right) = \left(\begin{array}{c} \text{Accuracy of the} \\ \text{Individual Forecast} \end{array} \right) \times \sqrt{n}$$

where n is the number of individual forecasts.

This says if we consolidate 25 locations the forecast is 5 times more accurate.

Let's return to our protection buffers. They are set to protect against the forecasted variability. (We are assuming future variability will look like past variability.) Placing assets at a number of forward distribution centers to provide protection buffers at each location requires more assets than placing the protection buffer at the source of supply (repair, production or purchasing) and using fast transportation to get them to the customers when needed.

How many more assets are we talking about? The Square Root Law says:

$$\left(\begin{array}{c} \text{The assets needed} \\ \text{for individual} \\ \text{protection buffers} \end{array} \right) = \left(\begin{array}{c} \text{The assets needed} \\ \text{for a centralized} \\ \text{protection buffer} \end{array} \right) \times \sqrt{\left(\begin{array}{c} \text{The number} \\ \text{of individual} \\ \text{protection buffers} \end{array} \right)}$$

If we choose to establish individual protection buffers at 25 locations for a given item, it will take 5 times more assets than would be needed to establish a consolidated protection buffer at the source of supply. For a given number of assets, a consolidated protection buffer would provide all the locations with 5 times more protection than would be possible if 25 individual protection buffers were established.

Sounds great, right? Well, before we sign up, let's examine the two assumptions on which the Square Root Law is based:

- Demands at each location are uncorrelated.
- The variability of demand is the same at all locations.

If the assumptions are valid, the Square Root Law holds. As the validity of the assumptions degrades, the benefit gained from centralizing the protection buffers decreases.

We are not talking about stripping all of the assets from the forward distribution centers and putting them in consolidated protection buffers. We are only talking about the assets needed to protect the customers from variability—safety stock. The requirement for assets to fill (the *average* portion of) the pipelines will still exist, although it should decrease as we use fast transportation to shorten the average length of our pipelines.

However, the Square Root Law clearly shows the benefit of consolidating the safety stocks, especially if it is done in combination with the creation and active management of protection buffers to allow us to execute based on actual real-time customer demand.

At this point, we start to see that *leaning* our logistics process requires us to implement an integrated solution. In addition, the consequences of failing to implement a piece of this integrated solution should be obvious. For example, we now understand that attempting to implement a demand-based logistics process without establishing and managing protection buffers does not solve the conflict represented in the Logician's Paradox.

Let's continue by examining each piece of our lean solution and how they fit together to form an integrated lean logistics process.

Establish Consolidated Protection Buffers

As we have already seen, it is essential that we establish protection buffers if we are to loosen ourselves from the grip of the Logician's Paradox. Prudence then directs us to consolidate those protection buffers whenever possible to benefit from the reduction of assets required to protect the customers. It is not enough to just establish protection buffers. It is essential that we continually measure the variability in our logistics process and adjust our protection buffers as needed to ensure our customers are protected from the effects of that variability.

Execute Based on Actual Customer Demands

With the establishment of protection buffers, we can safely move to demand-based execution of our logistics process. If the resources do not exist to satisfy all of the customers' requirements, then we prioritize those requirements and draw a cut line.

Once it is determined each day which requirements will be satisfied, each of those requirements becomes a demand on a protection buffer. If serviceable assets are available, they are immediately shipped to satisfy the customer requirements. If not, then the serviceable assets are shipped as soon as they

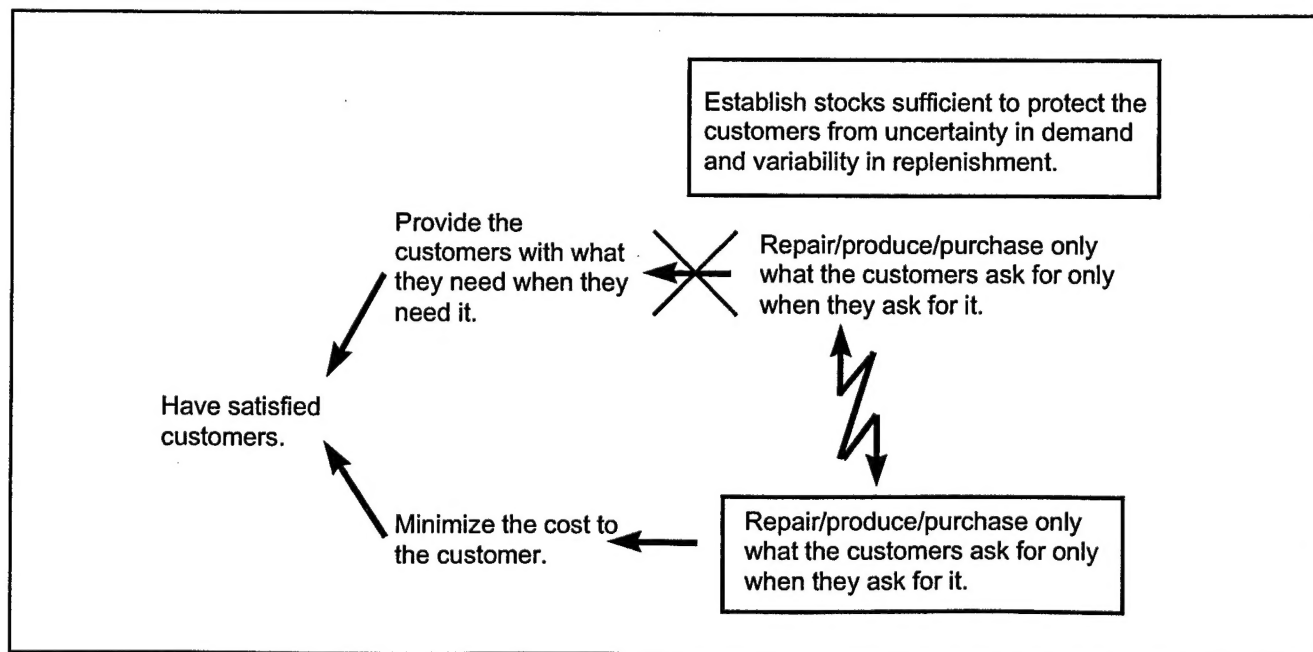


Figure 5. A Way Out of the Paradox

become available. To keep the flow of assets balanced, whatever number of customer requirements is to be satisfied, the same number of assets must be put into the logistics process. (The only exception would be if we made a change to the size of a protection buffer.) This is accomplished using a pull, pull, pull process. Here is how it works:

To those in the process upstream of the consolidated protection buffer, each demand placed on the consolidated protection buffer is viewed as a *hole*. Because they are graded on how long they take to fill each hole and by how much stuff is in their part of the pipeline, they work hard to make sure none of their protection buffers go *negative balance* (zero is acceptable in many cases) and to minimize the amount of stuff in their part of the pipeline.

When a demand is placed on the consolidated protection buffer, those responsible for *feeding it* reach back into their *assembly buffers* to get what they need to produce a serviceable end item which will satisfy the demand. Those components that were removed from an assembly buffer to produce the end item further pull on the logistics process by creating holes in the assembly buffers. Similarly, those responsible for feeding the assembly buffers reach back into their

parts buffers to get whatever they need to produce the components needed to fill the holes in their assembly buffers. In turn, this creates holes in the parts buffers, which pulls on those responsible for feeding the parts buffers.

This process of pull, pull, pull ensures the flow of assets is maintained. It synchronizes the efforts and resources of the entire logistics process and greatly simplifies the operating rules for those working in the logistics process.

Eliminate Queue or Wait Time

Take a process and ask, how long does it take to complete it—what is the *total process time* from start to finish. If we dissect the total process time, for nearly any process, we will find that 8 to 12 percent of the total process time is hands-on time. The rest is queue time—time where the work is just sitting and waiting. Now apply the Pareto principle and ask, where should we focus our process improvement efforts? Does it really make sense to spend money on new technology so we can put the paint on a widget in 3 minutes instead of 5, if those widgets sit for 3 1/2 weeks before they go into this paint booth?

An important piece of the lean solution is to squeeze the queue
(Continued on bottom of page 42)

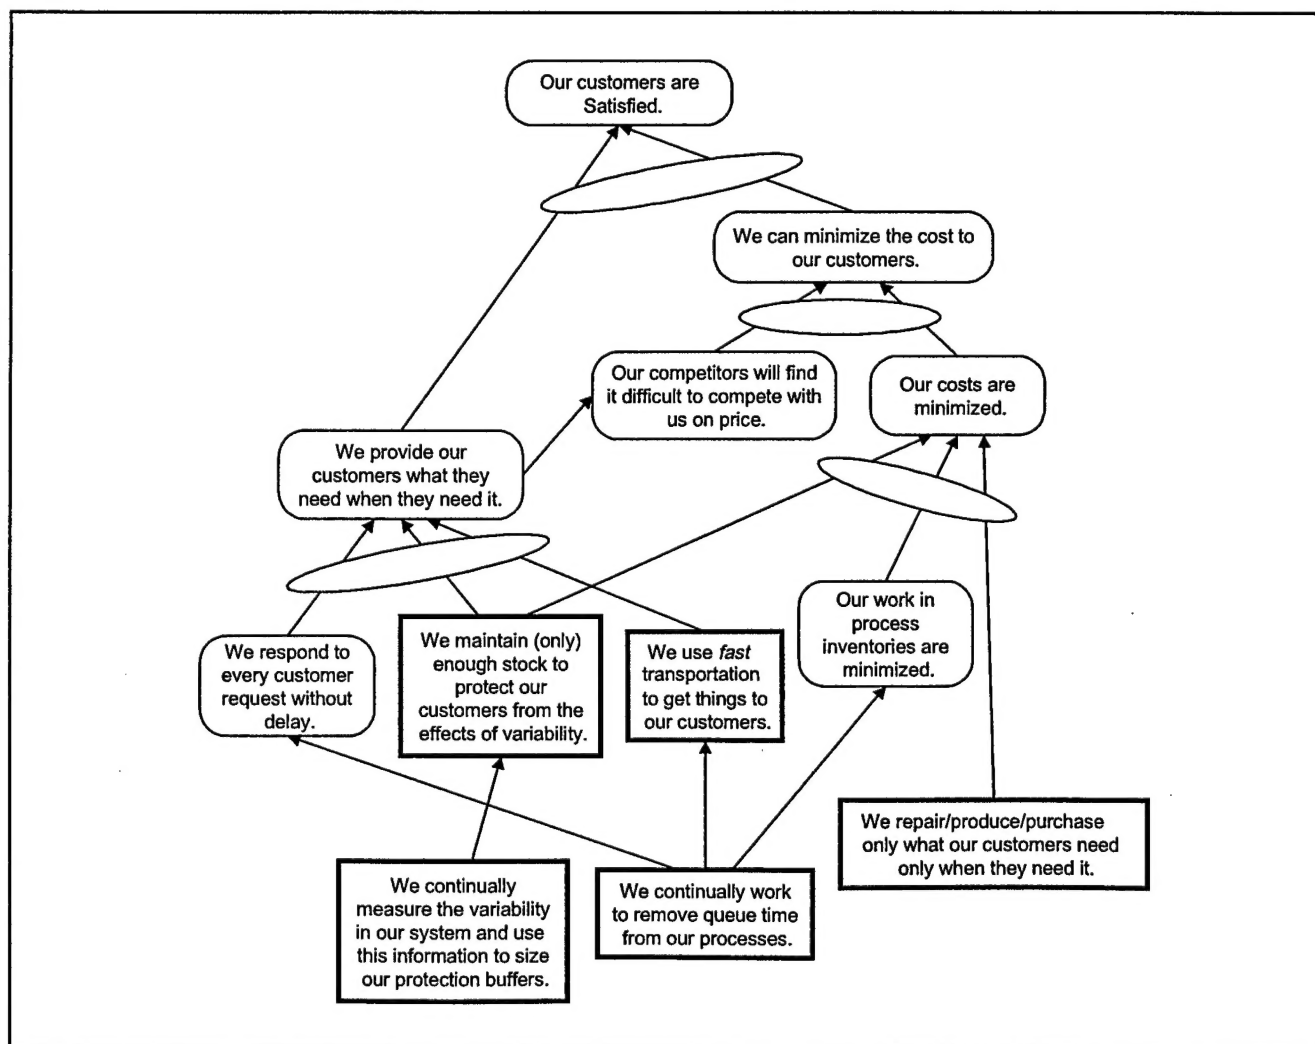


Figure 6. Putting All the Pieces Together to Form a Solution

time out of the process. Doing so reduces the average pipeline length and thus the amount of stuff needed to fill (the average portion of) the pipeline. Generally reducing the queue time also reduces the variability in the process. If the variability is reduced, then we do not need as much stuff in our protection buffers—another savings. As you can see, reducing queue time is very important.

Use Fast Transportation Everywhere

First, let's define what is meant by fast. When an asset is moved fast, it experiences little or no queue time. It does not wait for a cart, pallet, truck or whatever to fill up before it is moved to the next step in the process. As the engineers say, stuff is moved in *transfer batch* sizes of one. This is the most important piece of fast transportation. However, fast also means moving the asset in the quickest way practical. For most items, this means next-day air or dedicated truck.

Now that we have an understanding of each piece of the lean

logistics solution, let's take a look at Figure 6, which ties them all together.

Each of these bold square boxes contains a piece of the lean logistics solution. To understand how these pieces fit together to support the objective at the top, read each of the arrows in Figure 6 from tail to tip as *If . . . Then* statements, where the ellipses serve to indicate logical *ands*.

As Figure 6 highlights, in order to achieve the potential benefits of leaning our logistics processes, it is necessary to implement all the pieces. Recognizing and understanding the interrelationships between these pieces is key to successfully eliminating the age-old problem of logistics. It is much like baking a cake—none of the steps is overly difficult. However, if we skip a step or leave something out, the result usually falls far short of our expectations.

Norman Patnode is a first time contributor to the Air Force Journal of Logistics.